PCT/SE99/00605
09/673423
SE99/605

PRV

PATENT- OCH REGISTRERINGSVERKET
Patentavdelningen

REC'D 17 JUN 1999
WIPO PCT

Intyg
Certificate

Härmed intygas att bifogade kopior överensstämmer med de handlingar som ursprungligen ingivits till Patent- och registreringsverket i nedannämnda ansökan.

This is to certify that the annexed is a true copy of the documents as originally filed with the Patent- and Registration Office in connection with the following patent application.

(71) Sökande Net Insight AB, Stockholm SE
Applicant (s)

(21) Patentansökningsnummer 9801335-2
Patent application number

5

(86) Ingivningsdatum 1998-04-17
Date of filing

Stockholm, 1999-05-27

För Patent- och registreringsverket
For the Patent- and Registration Office

Emma Johnsson

Emma Johnsson

Avgift
Fee

PRIORITY DOCUMENT
SUBMITTED OR TRANSMITTED IN
COMPLIANCE WITH
RULE 17.1(a) OR (b)

PATENT- OCH REGISTRERINGSVERKET
SWEDEN

Postadress/Adress	Telefon/Phone	Telex	Telefax
Box 5055	+46 8 782 25 00	17978	+46 8 666 02 86
S-102 42 STOCKHOLM	Vx 08-782 25 00	PATOREG S	08-666 02 86

METHOD AND DEVICE FOR ALLOCATING TIME SLOTS TO A CHANNEL OF A CIRCUIT SWITCHED SYNCHRONOUS TIME DIVISION MULTIPLEXED NETWORK

Technical Field of Invention

The present invention relates to a method and a device for allocating time slots to a channel which is established to comprise one or more time slots in each recurrent cycle of a bitstream of a circuit switched synchronous time division multiplexed network.

Technical Background and Prior Art

Today, new types of circuit-switched communication networks are being developed for the transfer of information using circuit switched synchronous or isochronous, time division multiplexed bitstreams, wherein a bitstream is divided into cycles, each cycle in turn being divided into time slots.

One example of such a network is the so called DTM network (DTM - Dynamic Synchronous Transfer Mode). DTM is a broadband network architecture (see e.g. Christer Bohm, Per Lindgren, Lars Ramfelt, and Peter Sjödin, The DTM Gigabit Network, Journal of High Speed Networks, 3(2), 109-126, 1994, and Lars Gauffin, Lars Håkansson , and Björn Pehrson, Multi-gigabit networking based on DTM, Computer Networks and ISDN Systems, 24(2), 119-139, April 1992).

The time slots are preferably divided into two groups, control slots and data slots. Control slots are used for transferring of signaling messages between said nodes for the network's internal operation. The data slots are used for the transfer of data between said users connected to the different nodes.

Each node is arranged to dynamically allocate time slots for its respective end users to use when sending or receiving information to or from other users. Hence, different end users have access to different time slots.

Prior art describes rules for deciding, at a node, when the node shall give away time slots to other nodes. For example, as disclosed in WO 97/36402, a node will give away time slots if it has any free slots available, i.e. slots not allocated to any one of the end users served by said node. However, in many cases, there will be a need for more flexible rules, providing more options in the control of allocation and deallocation of time slots. Also, a problem will arise if there are no free slots available on the bitstream of interest.

An object of the invention is hence to provide a more flexible, yet simple, way of controlling the allocation and deallocation of time slots to channels, also taking into consideration for example the situation when there are no free slots available.

Summary of the Invention

This and other objects are accomplished by the invention as defined in the accompanying claims.

Hence, according to a first and a second aspect of the invention, there is provided a method and a device, respectively, of the kind mentioned in the introduction, wherein a time slot allocated to said channel is associated with a selected level, of at least two available levels, of priority by which said channel owns said time slot.

The invention is thus based upon the idea of assigning a selected level of priority, also called prioritized ownership, to the bond between a channel and a time slot allocated to said channel. The decisions as to whether or not to deallocate said time slot from said channel, generally as a result of a request for bandwidth to be allocated to another channel, may then be based upon the assigned level of priority by which the channels own said slot.

Typically, the decision as to whether or not a time slot shall be deallocated from said channel will be based

upon a level of priority associated with a request for a time slot for another channel. If a channel A uses less prioritized slots, another channel B that needs more resources, which cannot be satisfied by free slots, can force deallocation of the "softly" allocated slots from channel A for use in channel B if its request refers to a more prioritized ownership.

As is understood, a request of this kind may be both generated as well as accommodated at one single node, the accommodation of said request thus not requiring any sending of requests to other nodes. However, the accommodation of said request may of course involve the sending of the request to other nodes. In the latter case, according to an embodiment, the invention of course not being limited thereto, only free slots is reallocated among nodes to accommodate the request if the request is sent to another node, i.e. the priority associated with a request is ignored in such a case, the priority associated with said request only being considered when allocating slots between channels handled by the one single node.

The invention thus allows for the use of resources allocated to a circuit-switched channel in a more flexible way. If, for example, there are free resources existing on a bitstream, i.e. free, non-allocated time slots, these may be used temporarily by a first channel until a second channel with a higher prioritized need for these resources appears. The resources will then be deallocated from the first channel and allocated to the second channel. Resources may thus by used by merely best-effort requiring channels (circuits) without the risk of blocking higher prioritized channels. The scheme therefore provides a possibility to use the resources more efficiently in the case of best-effort traffic, such as data traffic, in a circuit-switched network.

An advantage of the invention is that it provides a scheme which will allow, for example, a network operator

to offer different user service classes to its end user or customers. For example, a high service class, giving access to channels owning slots with a high level of priority, and consequently providing a more reliable access to time slots, may be offered to customers having such needs, such as a television broadcasting company, whereas a lower service class, giving access to channels owning slots with a lower level of priority, and consequently providing a less reliable access to time slots, may be offered to customers having such limited needs.

As is understood, in alternative embodiments, there may be more than two levels of priority. Also, each level of priority, as well as each type of service class, may be associated with different kinds of characteristics. For example, the channel set up characteristics of a certain priority level may differ from the deallocation characteristics of said level. As an example, the characteristics of a first level may be such that a request referring to said first level will not have priority over a simultaneously provided request referring to a second level, even though a slot already allocated to a channel at said first level of priority will not be deallocated as a result of a request referring to said second level of priority. Preferably, the characteristics of the different levels of priority may thus be selected and decided by the network operator and may even be customized based upon the expressed needs of the customer or end user.

The selected level of priority by which a channel owns a time slot is preferably decided at channel set up. After the selected level of priority has been assigned to the slot it may, however, be subject to change during the life time of the channel as a consequence of changing bandwidth requirements. For instance, an end user requesting a channel may have low requirements as to when bandwidth is allocated to said channel, but may have high requirements as to the importance of an uninterrupted

access once the requested bandwidth is allocated. This would be accomplished by the sending of a request for bandwidth, said request being associated with a low level of priority, and by allocating, once the request has been met, time slots to said channel and assigning these time slots with a high level of priority.

A channel may very well have different levels of prioritized ownership assigned to different slots that it owns, or it may have the same level assigned to all its slots. If, for example, a circuit is set up with half of the slots using the highest available prioritization and the rest using a lower prioritization, the circuit is always ensured it will keep a certain capacity, corresponding to the number of highly prioritized slots.

Also, time slots of a channel may be assigned different levels of priority at different parts of the network. For example, in a tree-like hierarchical bitstream structure, this may be used to ensure that the slots of the channel will be owned at a high priority in the larger trunks of the tree but at a lower priority at the leafs of the tree. For example, this feature will provide a very advantageous tool in the management of multicasted channels.

As is understood, there may be different strategies for deallocating time slots from channels having time slots all owned at the same priority level when there is an existing bandwidth shortage. For example, such strategies may follow rules such as: the last time slot allocated to any channel (at the given priority) will be deallocated first (i.e. stack type of strategy); the longest (in time) allocated time slot will be deallocated first (i.e. a FIFO type of strategy); deallocating slot in a round robin fashion (i.e. deallocate one time slot at a time from each channel); or deallocation in a manner which provides a best fit considering the time slot fragmentation on the bitstream of interest. Of course, these strategies will basically be no different than the ones

that may be used in prior art, wherein slots are either unallocated, i.e. free, or allocated (without any choice of allocation strength).

Also, since the invention relates to circuit switched networks, the priority assignment according to the invention does basically not refer to the priority by which a time slot is to be passed along through the network, i.e. the kind of priority used in packet switched networks and the like for ensuring priority in case of a blocking situation occurring somewhere along the packet travel path through the network. Instead, the priority levels according to the invention are used primarily in the management of channel set-up and of channel bandwidth changes.

Although it has been stated above that a priority level according to the invention refers to the bond between a time slot and the channel to which said slot is allocated, it is to be understood that any assignment of a priority level explicitly referring to, for example, the bond between a time slot and an end user, between a time slot and a node, between a channel and a node, or the like, will inevitably involve or imply, directly or indirectly, explicitly or implicitly, the provision of a priority level referring to the bond between a time slot and a channel.

Further aspects, objects and features of the invention will be understood more fully from the accompanying claims and from the following description of exemplifying embodiments thereof with reference to the accompanying drawings.

Brief description of the drawings

Exemplifying embodiments of the invention will now be described with reference to the accompanying drawings, wherein:

Fig. 1 schematically shows an exemplified network in accordance with the invention;

Fig. 2 schematically shows an exemplifying time multiplexed bitstream propagating along one of the optical fibers shown in Fig 1;

Fig. 3 schematically shows a block diagram of a node according to an embodiment of a device for allocating time slots according to the invention;

Fig. 4a shows an exemplified time slot allocation diagram according to an embodiment of the invention;

Fig. 4b shows the time slot distribution among prioritized ownership levels in accordance with the allocation shown in Fig. 4a;

Fig. 5a and 5b show flow diagrams according to an embodiment of the invention; and

Figs. 6a to 6d show exemplifying slot utilization tables according to another embodiment of the invention.

Detailed description of preferred embodiments

In Fig. 1, the basic topology of a synchronous time multiplexed network according to the invention is shown. The network in Fig. 1 comprises three nodes N1, N2 and N3, each connected to a bus having two unidirectional optical fibers B1 and B2 connecting all nodes. The optical fiber B1 carries at least one bitstream used for communication in one direction along the bus, and the optical fiber B2 carries at least one bitstream used for communication in the other direction along the bus, as indicated in Fig. 1 by the arrows at the end of the optical fibers B1 and B2. Even though a simple two-way structure is shown in Fig. 1, a ring structure or a hub structure or the like may also be used. Also, two or more busses (or for that matter two or more single fibers) may be interconnected using switch nodes, thereby creating a network having the structure of a two dimensional or three dimensional mesh. As will be described below, the communication on each optical fiber B1 and B2 is synchronous and time multiplexed. In addition, wavelength division multiplexing, i.e. having each of the fibers carry

more than one bitstream using different wavelengths, may be used to increase the network capacity.

As shown in Fig. 1, each node is arranged to serve one or more end users by providing access to the optical fibers B1 and B2. Hence, three end users 11, 12 and 13 are connected to the first node N1, three end users 21, 22 and 23 are connected to the second node N2, and three end users 31, 32 and 33 are connected to the third node N3. When, for example, the end user 21 connected to the second node N2 wants to send information to the end user 32 connected to the third node N3, the nodes N2 and N3 will establish a communication channel on a bitstream propagating on the optical fiber B1. The second node N2 will then transfer data from the sending end user 21 to the bitstream on the optical fiber B1, and the third node N3 will in turn transfer said data from the bitstream on the optical fiber B1 to the receiving end user 32. If information is to be passed from the end user 32 to the end user 21 as well, a similar channel is set up on a bitstream propagating on the optical fiber B2. Even though the end users have been shown as computers or similar work stations in Fig. 1, it is understood that the end user may be any type of electronic equipment needing access to the network, such as printers, servers, facsimile machines, telephones, television sets, radio receivers, and the like.

The structure of a time multiplexed bitstream of the kind propagating on the optical fibers B1 and B2 in Fig. 1 will now be described with reference to Fig. 2. The bandwidth of each wavelength, i.e. each bitstream, is divided into 125 µs cycles. Each cycle is in turn divided into 64 bit time slots. The number of time slots within a cycle thus depends on the network's bit rate. Consequently, the number of slots shown in the cycle of the bitstream in Fig. 2 is merely illustrative, the actual number of slots within each cycle being far greater than what is shown in Fig. 2.

The time slots are in general divided into two groups, control slots C and data slots D. The control slots C are used for control signaling between nodes in the network, i.e. for carrying messages between nodes for the internal operation of the network, such as for channel establishment, slot allocation, and the like. The data slots D are used for the transfer of user data between end users connected to said nodes.

In addition to said control slots and data slots, each cycle comprises one or more synchronization slots S used to synchronize the operation of each node in relation to each cycle or frame. Also, a guard band G is added after the last slot at the end of each cycle in order to facilitate synchronization. As indicated in Fig. 2, the bitstream cycle is repeated continuously.

Each node has access to at least one control slot C and to a dynamic number of data slots D. Each node uses its control slot C to communicate with other nodes within the network. The number of data slots D to which a node has access depends upon the transfer capacity requested by the end users served by the respective node. However, a node may also have access to data slots which are not needed to satisfy the transfer capacity requested by the users, or channels. These slots are then denoted non-allocated, or free, slots. If the end users of a certain node require a large transfer capacity, the node will allocate more data slots for that purpose, which may involve reallocation of free time slots among different nodes. On the other hand, if the end users of a certain node merely requires a small transfer capacity, the node may limit its number of allocated data slots. However, a node may continue to have access to data slots which at the moment are not needed to satisfy the transfer capacity currently requested by the users. Also, the number of control slots to which each node has access to may be increased or decreased depending on the node's demand for signaling capacity. Hence, the number of data slots as

well as control slots to which a node has access to may by dynamically adjusted as the network load changes.

In Fig. 2, the first node N1 has access to one control slot and one data slot. The second node N2 has access to one control slot and three data slots as a result of its end users currently having a greater need to send data. The third node N3 has merely access to a control slot and no data slots as a result of its users temporarily having no need for sending data via the network or currently being arranged to merely receive data.

With reference to Fig. 3, an embodiment of a device for allocating time slots to a channel in accordance with the present invention will now be described. Fig. 3 shows the basic components of an exemplified network node in a synchronous time division multiplexed network. This node could for example be any one of the nodes N1, N2 and N3 in Fig. 1. The node 100 comprises a synchronization detector 110, a time slot counter 120, a logic unit 130, a network interface 140, a user interface 150, a memory 160 and a node controller 170. The node controller 170 comprises a unit 180 for allocating time slots to a channel and a slot utilization table 176. The memory 160 comprises a write data table 162 and a read data table 164. The user interface 150 is connected to the end users attached to the node 100, such as the end users 11, 12 and 13, or the end users 21, 22, and 23 in Fig. 1.

The synchronization detector 110 will derive the cycle or frame synchronization signal (designated S in Fig. 2) from the bitstream passing on an optical fiber (not shown) connected to the node 100 and use it to restart the time slot counter 120. The counter 120 will count the number of time slots passing on the optical fiber at a predefined bit rate and provide a corresponding signal to the logic unit 130 and the network interface 140. The logic unit 130 uses the signal from the counter 120 to keep track of which time slot within the

cycle that is being currently processed, and the network interface 140 uses the signal from the counter to synchronize the writing into, or the reading of, slots with the passing bitstream.

The node controller 170 will keep track of all required information as to the allocation of slots to different nodes and different channels. Hence, the node controller is for example used when establishing new channels on behalf of the end users connected to the user interface 150. The node controller has specified in the write data table 162 of memory 170 in which slots the node may write control data and user data. In the read data table 164 of memory 170, the node controller has specified from which slots the node may read control data and user data. Control data is transferred to the node controller and user data is transferred to the appropriate user connected to the node. In the slot utilization table 176, the unit 180 for allocating slots to a channel has specified the level of prioritized ownership by which a channel owns the respective time slot. Different slots may be owned by different channels, or by the same channel, with different levels of prioritized ownership, ranging from a highest priority level to a lowest priority level. In Fig. 3 it is specified in table 176 that slot 3 is owned by channel A with a priority level 2. Slot 5 is also owned by channel A but with a priority level 3 and slot 9 is owned by a channel C also with a priority level 3.

The unit 180 comprises priority assignment means 182 and slot allocating means 184. The priority assignment means assigns to a time slot owned by (or to be owned by) a channel a level of prioritized ownership by means of writing information designating this level in the slot utilization table. The assigned level corresponds to a level selected by an end user when the end user makes a request for either its initial bandwidth needs or for its changed bandwidth needs. The term "bandwidth needs" may

in this context refer to the size of the bandwidth or the priority level of the bandwidth, or both. The slot allocating means 184 are arranged to request a time slot for a channel in need of bandwidth and to include a user selected level of prioritized ownership in this request. The slot allocating means 184 are also arranged to receive any such a request from an end user connected to this node or from another node on behalf of an end user connected thereto.

An example of the allocation and deallocation of time slots to channels, or end users, according to one embodiment of the invention will now be described with reference to Figs. 4a and 4b. In this embodiment, a channel may own, or have access to, a time slot at one of three different possible levels of priority. The three different levels have the following characteristics:

The highest level of priority is a reserved level, denoted R in Fig 4b. A time slot owned by a channel at this level of priority will be called a reserved time slot. A reserved time slot may not, in this exemplifying embodiment, be deallocated as a result of a request for transfer capacity even if the channel is not using the time slot for transferring data. That is, a reserved time slot which is not in use will not become a free time slot. As an example, a leased line would comprise a number of reserved time slots.

The second highest level is a so called hard level, denoted H in Fig 4b. A time slot owned by a channel at this level of priority will be called a hard time slot. A hard time slot may, in this exemplifying embodiment, only be deallocated as a result of a request for a reserved time slot. Hence, neither a request for a soft time slot, to be described below, nor a request for a hard time slot can force a hard time slot to be deallocated. When a channel stops sending data in a hard time slot, the hard time slot is preferably deallocated and it becomes a free slot.

The lowest level of priority is called a soft level, denoted S in Fig 4b. A time slot owned by a channel at this level of priority will be called a soft time slot. A soft time slot may, in this exemplifying embodiment, be deallocated as a result of a request for a reserved time slot or for a hard time slot. A request for a soft time slot can however not force a soft time slot to be deallocated. When a channel stops sending data in a soft time slot, the soft time slot is preferably deallocated and becomes a free slot.

It should be noted that an alternative way to implement the highest level of ownership, that is the reserved level, is to let reserved time slots, which currently are not used by its owner for transferring data, temporarily be owned by another channel with a lower level of ownership, in which case the time slot is immediately deallocated from this other channel as soon as the channel to which the time slot is reserved wants to use the slot.

Fig. 4a illustrates an exemplified allocation of 15 time slots in a cycle of a bitstream. As discussed above, in Fig. 4a, a time slot is either free F or owned/used by any one of the channels A, B, C or D. A reserved time slot not being used by its owner is indicated NU (Not Used). In Fig. 4b the slot distribution among the different levels of prioritized ownership is shown. Hence, Fig. 4b shows whether or not a time slot in the cycle is free F or owned by a channel, and if owned by a channel, the level of prioritized ownership, i.e. a reserved level R, a hard level H, a soft level S.

In this example, it is assumed that channel A will request 2 reserved time slots, that channel B will request 2 hard and 8 soft slots, that channel C will request 4 hard slots, and that channel D will request 4 hard and 4 soft slots.

It should be noted that the time axis used in Figs. 4a and 4b are merely for illustrating the relative order

of certain actions affecting the allocation of time slots in the cycle.

At time t=0, channel A requests 2 reserved time slots, leaving the cycle with 13 free time slots. The transfer capacity of these two reserved slots is initially not used by channel A, as is indicated in the slot utilization of Fig. 4a.

At t=2, channel B requests 2 hard time slots and 8 soft time slots. Since there are 13 free slots in the cycle, 2 hard slots and 8 soft slots will be allocated to channel B to be owned by channel B in accordance with its request. In Fig. 4a it is seen that 10 slots are now utilized by channel B, and in Fig. 4b it is seen that the cycle now consists of 2 reserved slots, 2 hard slots, 8 soft slots and 3 free slots.

At t=4, channel C requests 4 hard time slots. The three free slots is therefore allocated to channel C as hard slots. One of the soft slots allocated to channel B is now deallocated, since these slots is owned by channel B at a lower level of prioritized ownership as compared to the priority referred to by the request from channel C regarding a hard slot. The soft slot deallocated from channel B is allocated as a hard slot to channel C. In Fig. 4a it is seen that 9 slots are now utilized by channel B and 4 slots by channel C, and in Fig. 4b it is seen that the cycle now consists of 2 reserved, 6 hard and 7 soft slots.

At t=6, channel A starts using its transfer capacity which has previously been reserved. At t=8, all slots allocated to channel B are deallocated and become free slots, since the bandwidth need of channel B temporarily ends at this time. At t=10, channel D makes its bandwidth request for 4 hard slots and 4 soft slots, and 8 of the 9 free slots are therefore allocated to channel D.

At t=12, channel B ones again makes a request for 2 hard and 8 soft slots. The one remaining free slot is then allocated to B as a hard slot and one of the soft

slots allocated to D is deallocated and allocated to B as the second hard slot. However, channel B's request for 8 soft slots cannot be met, and this part of the request has to be rejected.

At t=14, channel A stops using its reserved slots for transferring data. As previously described these slots are however not deallocated but continue to be reserved for channel A. At t=17, channel C stops using its transfer capacity and its 4 hard slots are therefore deallocated to become free slots.

At t=19 channel B again tries to gain access to the whole size of its needed bandwidth. Since 2 hard slots are already allocated to channel B, the request will be for the additional 8 soft slots. At this point the 4 free slots will be allocated to B as soft slots while the request for the remaining 4 soft slots will still be rejected.

It is to be noted that the example discussed above may involve reallocation of time slots between nodes (so called inter node allocation) as well as within a node (so called intra node allocation).

Flow diagrams of processing steps performed in the managing of allocation of time slots to a channel in accordance with the present invention will now be described with reference to Figs. 5a and 5b.

In Fig. 5a, the processing starts in step 510 with the reception of a request for a slot to be allocated to a channel at an assigned selected level of ownership. In Fig. 5a, it is assumed that the request comes from an end user connected to the node performing the illustrated steps.

In step 520, the node checks if there are any free slots available in the slot utilization table of this node. If the answer is yes, the process continues to step 580 followed by step 590. In step 580, such free slots are allocated to the requesting channel, and in the following step 590, the selected level of prioritized

ownership is assigned to the slot allocated in step 580. If however the answer in step 520 is no, the process continuos to step 530.

In step 530, the slot utilization table of the node is examined for any slots having a lower level of priority than the selected level associated with the request received in step 510. If any such slots are found, the process continuos to step 570. If no such slots exists, the process continuos to step 540. In step 570, the slots found in step 530 are deallocated from the channel or channels being the current owner of the slots. These deallocated slots are then processed in step 580 and 590 as previously described. In step 540, the slot request is forwarded to other nodes in association with the selected level of priority for the desired slot.

In step 550, answers received as a result of the request in step 540 are then examined. The answers will indicate whether or not the access rights to any slots may been transferred from other nodes to this node. If the answer indicates that the access right to a slot has been transferred, steps 580 and 590 described above is performed with respect to this slot. If no access right to a slot has been transferred, step 560 is executed. In step 560, the request for a slot with the selected level of priority is rejected.

In Fig. 5b, the processing steps refer to a request for a slot which should be assigned a selected level of ownership, said request being derived from an end user connected to a different node than the one performing the illustrated steps, said node having forwarded the request by means of executing the previously described step 540 in Fig. 5a.

The processing starts in step 610 with the reception of the request. After reception of the forwarded slot request, it is checked in step 620 whether or not there are any free slots in the slot utilization table of this

node. If the answer is yes, step 650 followed by step 680 are executed. Otherwise, step 630 is executed.

In step 650, the node will initiate the transfer of the access right to the free slots to the node from which the request was received in step 610. In step 680, the transferred slot will be removed from the slot utilization table. In step 630, the slot utilization table of the node is examined for any slots having a lower level of priority than the requested level. If such a slot is found, step 660 is processed. If no such slot exists, step 640 is processed.

In step 660, a slot found in block 630 is deallocated from the channel being the current owner of the slot. This deallocated slot is then processed in step 670 followed by step 680. In step 670, the node will initiate the transfer of the access rights to the slot deallocated in block 660 to the node from which the request in block 610 was received, and in step 680, the transferred slot will be removed from the slot utilization table. In step 640, having found no deallocatable slots given the priority level associated with the request, the request received in step 610 will be rejected.

Figs. 6a to 6d show slot utilization tables, in this example relating to the slot utilization described with reference to Fig. 4a. These slot utilization tables are in this example of the kind previously described with reference to Fig. 3. It is assumed that Figs. 6a and 6c show the slot utilization table of node N1 in Fig. 1 at two different points in time and that Figs. 6b and 6d show the slot utilization table of node N2 in Fig. 1 at the corresponding points in time. It is also assumed that channel A is to be established from node N2 to node N3 and that channel B, C and D are to be established from node N1 to node N3.

At a time corresponding to t=3 of the diagram described in Fig. 4a, the nodes N1 and N2 have the slot utilization tables seen in Figs. 6a and 6b, respectively.

At t=4 in Fig. 4a, channel C requests 4 hard slots. In Fig. 6a, it is seen that node N1 prior to this request only has access to one free slot and that node N2 has access to two free slots.

Figs. 6a to 6d illustrate the consequences of the processing steps performed by the node controllers in node N1 and N2. When the node controller of node N1 receives the request for 4 hard slots from channel C at time t=4, it first allocates any free slots it has access to, in this case the single slot 13, to the channel and assigns the selected level of ownership received in the request to these slots. This leaves another 3 hard slots to be allocated to meet the request. The node controller in node N1 then forwards a request for 3 hard slots to node N2. Since node N2 has two free slots, slots 14 and 15, it will as a result of the request transfer the access to these two slots to node N1. The node controller of node N1 will then allocate these two slots as hard slots to channel C. Node N1 will then examine its slot utilization table for a slot owned with a lower ownership level than the requested ownership level to fulfill the request of channel C, find one of the soft slots allocated to channel B, for example slot 12, deallocate this slot from channel B, and allocate it to channel C as a hard slot.

Hence, this scheme utilizes all free slots, regardless of what nodes that have access to those free slots, before a slot owned with a level of prioritized ownership is deallocated on behalf of a request for a slot to be owned with a higher level of ownership, which differs somewhat as compared to the process described above with reference to Figs. 5a and 5b.

As described above with reference to Figs. 4a and 4b, the three free slots in the cycle will be allocated to channel C as hard slots and one soft slot allocated to channel B will be deallocated and then allocated to channel C as a hard slot. After the request from channel

C has been met, at time t=5, the nodes N1 and N2 will have the slot utilization tables shown in Figs. 6c and 6d, respectively.

Even though the invention has been described with reference to exemplifying embodiments thereof, these are not to be considered as limiting the scope thereof, and, as understood by those skilled in the art, different modifications and alterations may be made within the scope of the invention, which is defined by the accompanying claims.

CLAIMS

1. A method for allocating time slots to a channel which is established to comprise one or more time slots in each recurrent cycle of a bitstream of a circuit switched synchronous time division multiplexed network, said method comprising the step of associating a time slot allocated to said channel with a selected level, of at least two available levels, of priority by which said channel owns said time slot.

2. A method as claimed in claim 1, comprising the step of determining whether or not a time slot of said channel, said time slot being owned by said channel with a selected level, of at least two available levels, of priority, shall be deallocated from said channel based upon the level of priority associated with a request for a time slot for another channel.

3. A method as claimed in claim 1 or 2, comprising the step of deallocating a time slot of said channel, said time slot being owned by said channel with a selected level, of at least two available levels, of priority, as a result of a request for a time slot to be owned by another channel, said request referring to a higher level of priorit than said selected level of priority by which said channel owns said time slot.

4. A method as claimed in claim 1, 2, or 3, comprising the step of deallocating a time slot of said channel, said time slot being owned by said channel with a selected level, of at least two available levels, of priority, as a result of a request for a time slot to be owned by another channel with a higher level of priority than said selected level of priority by which said channel owns said time slot.

5. A method as claimed in claim 2, 3, or 4, wherein said step of deallocating a time slot of said channel is performed only if there are no non-allocated slots available.

6. A method as claimed in any one of claims 2 to 5, wherein said step of deallocating a time slot of said channel is performed only if said time slot is owned by said channel with a level of priority that is lower than a highest level of priority.

7. A method as claimed in any one of claims 2 to 6, wherein the decision to deallocate a time slot of said channel is further based upon an evaluation regarding to which channel a time slot was last allocated.

8. A method as claimed in any one of claims 2 to 7, wherein the decision to deallocate a time slot of said channel is further based upon an evaluation regarding to which channel a time slot has been allocated the longest period of time.

9. A method as claimed in any one of claims 2 to 8, wherein the decision to deallocate a time slot of said channel is further based upon an evaluation regarding from which channel a time slot was last deallocated.

10. A method as claimed in any one of claims 2 to 9, wherein the decision to deallocate a time slot of said channel is further based upon an evaluation regarding from which channel a time slot should be deallocated in order to counteract time slot fragmentation on the bitstream of interest.

11. A method as claimed in any one of the preceding claims, comprising the step of associating all time slots owned by said channel with the same selected level of

priority, said time slots thus being owned by said channel with the same selected level of priority.

12. A method as claimed in claim 11, wherein said step of associating all time slots owned by said channel with the same selected level of priority comprises associating said channel with said selected level of priority, each time slot allocated thereto thereby being owned by said channel with said selected level of priority.

13. A method as claimed in any one of claims 1 to 10, comprising the step of associating different time slots owned by said channel with different levels of priority, said time slots thus being owned by said channel with different levels of priority.

14. A method as claimed in any one of claims 1 to 10, comprising the step of associating time slots owned by said channel at one bitstream or bitstream section with a selected level of priority and associating slots owned by said channel at another bitstream or bitstream section with another selected level of priority.

15. A method as claimed in any one of the preceding claims, comprising the step of changing the level of priority by which a time slot is owned by said channel.

16. A method as claimed in any one of the preceding claims, comprising the step of dynamically allocating, or deallocating, time slots of at least one level of priority to, or from, said channel in accordance with the dynamically changing bandwidth needs of said channel.

17. A method as claimed in any one of the preceding claims, comprising the steps of:

receiving a request for a time slot, said request being associated with a requested level of priority; and

deallocating, if existing, a time slot owned by a channel with a lower level of priority than said requested level of priority for placing said time slot at the disposal of said request.

18. A method for allocating time slots to a channel which is established to comprise one or more time slots in each recurrent cycle of a bitstream of a circuit switched synchronous time division multiplexed network, comprising the steps of:

requesting a time slot for a channel in need of bandwidth, said request being associated with a level of priority;

allocating to said channel a time slot put at said channel's disposal as a result of said request; and

associating said time slot allocated to said channel with a selected level of priority, said channel thus owning said time slot with said selected level of priority.

19. A method as claimed in claim 18, wherein said selected level of priority is different than the level of priority associated with said request.

20. A method as claimed in claim 18, wherein said selected level of priority is not the same level as the one associated with said request.

21. A device for allocating time slots to a channel which is established to comprise one or more time slots in each recurrent cycle of a bitstream of a circuit switched synchronous time division multiplexed network, said device comprising:

priority assignment means (182) for associating a time slot allocated to said channel with a selected level, of at least two available levels, of priority by which said channel owns said time slot.

22. A device as claimed in claim 21, comprising a slot utilization table (176) indicating the selected level of priority by which said channel owns said time slot.

23. A device as claimed in claim 22, wherein said priority assignment means is arranged to write information designating the selected level of priority by which said channel owns said time slot into said slot utilization table.

24. A device as claimed in any one of claims 21 to 23, comprising slot allocating means (184) provided to request a time slot for a channel in need of bandwidth, said request referring to a selected level of priority; and to allocate to said channel a time slot put at disposal as a result of said request.

25. A device as claimed in any one of claims 21 to 24, comprising slot allocating means (184) provided to receive a request for a time slot and to deallocate a time slot of said channel, said time slot being owned by said channel with a selected level, of at least two available levels, of priority, based upon whether or not said request refers to a higher level of priority ownership than said selected level of priority by which said channel owns said time slot.

26. A device as claimed in any one of claims 21 to 25, wherein said device manages the slot allocation/-deallocation on behalf of at least one node of said network.

27. A device as claimed in any one of claims 21 to 26, wherein said device manages slot allocation/-deallocation on behalf of several nodes of said network.

28. The use of a method as claimed in any one of claims 1 to 20, or a device as claimed in any one of claims 21 to 27, for specifying different traffic service classes when operating said network.

ABSTRACT

The present invention relates to a method and a device for allocating time slots to a channel which is established to comprise one or more time slots in each recurrent cycle of a bitstream of a circuit switched synchronous time division multiplexed network. According to the invention, a time slot allocated to said channel is associated with a selected level, of at least two available levels, of priority by which said channel owns said time slot.

Elected for publication: Fig. 3

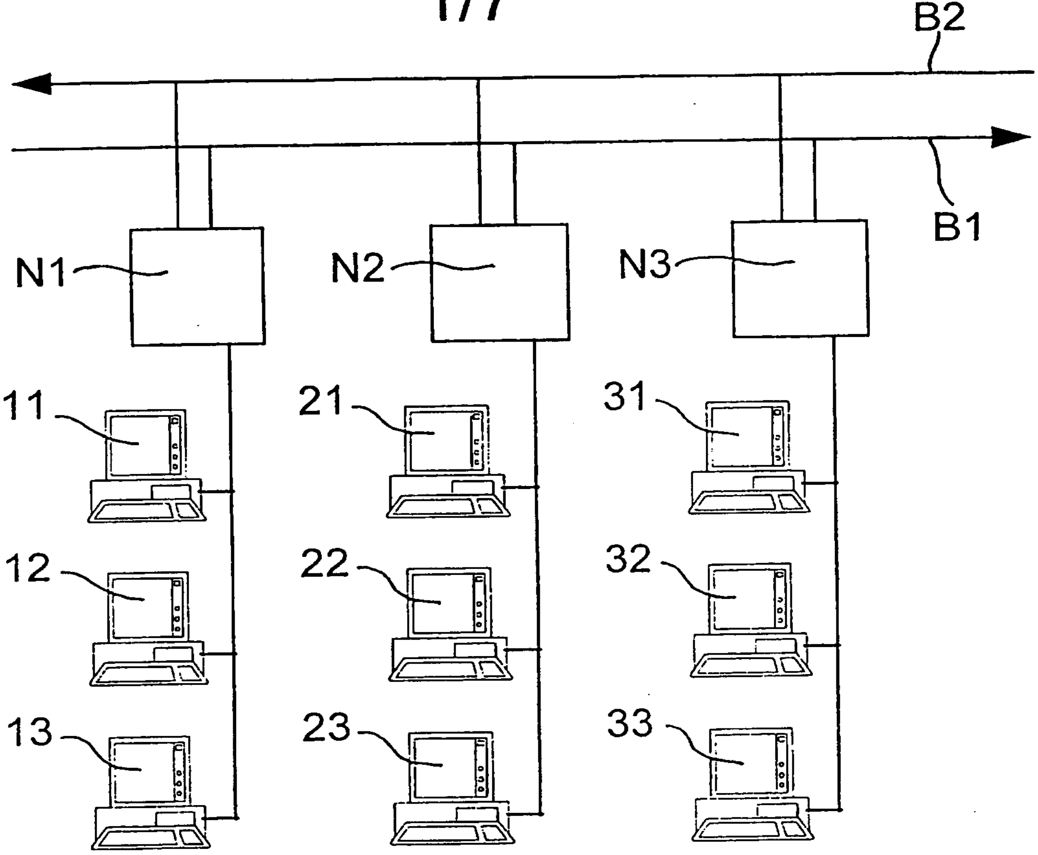
B2
B1
N1
N2
N3
11
12
13
21
22
23
31
32
33

FIG. 1

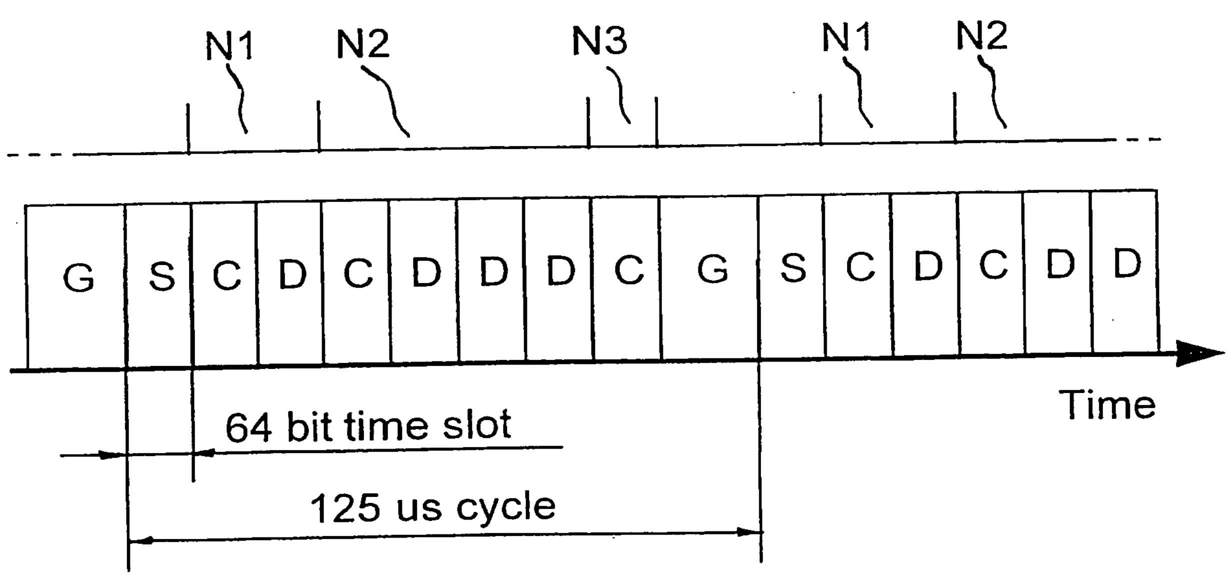
N1
N2
N3
N1
N2
G
S
C
D
C
D
D
D
C
G
S
C
D
C
D
D
Time
64 bit time slot
125 us cycle

FIG. 2

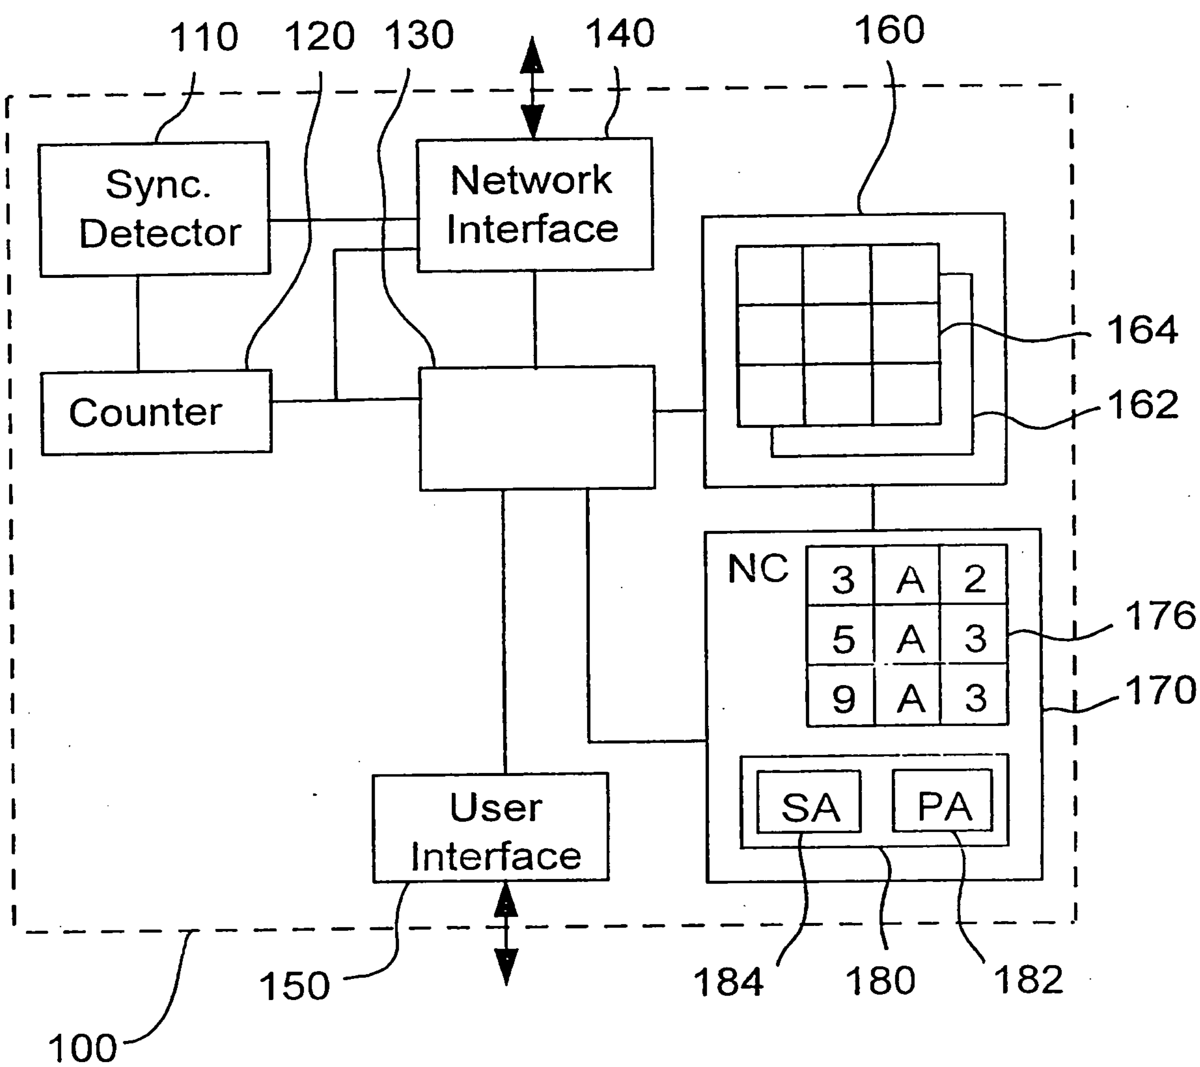
110
120
130
140
160
Sync. Detector
Network Interface
Counter
164
162
NC
3 A 2
5 A 3
9 A 3
176
170
User Interface
SA
PA
150
100
184
180
182

FIG. 3

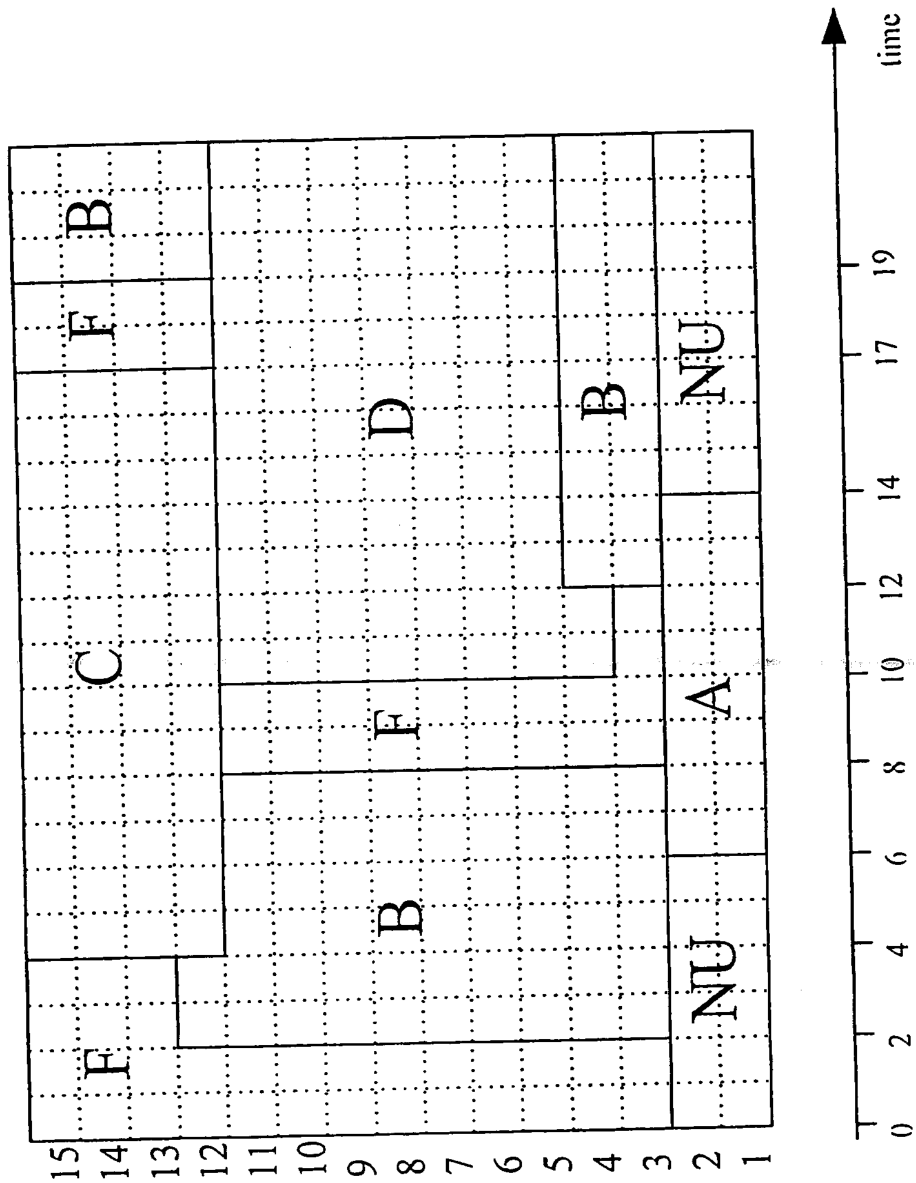

F
B
C
F
B
F
D
B
NU
A
NU
time
0
2
4
6
8
10
12
14
17
19
1
2
3
4
5
6
7
8
9
10
11
12
13
14
15

FIG. 4a

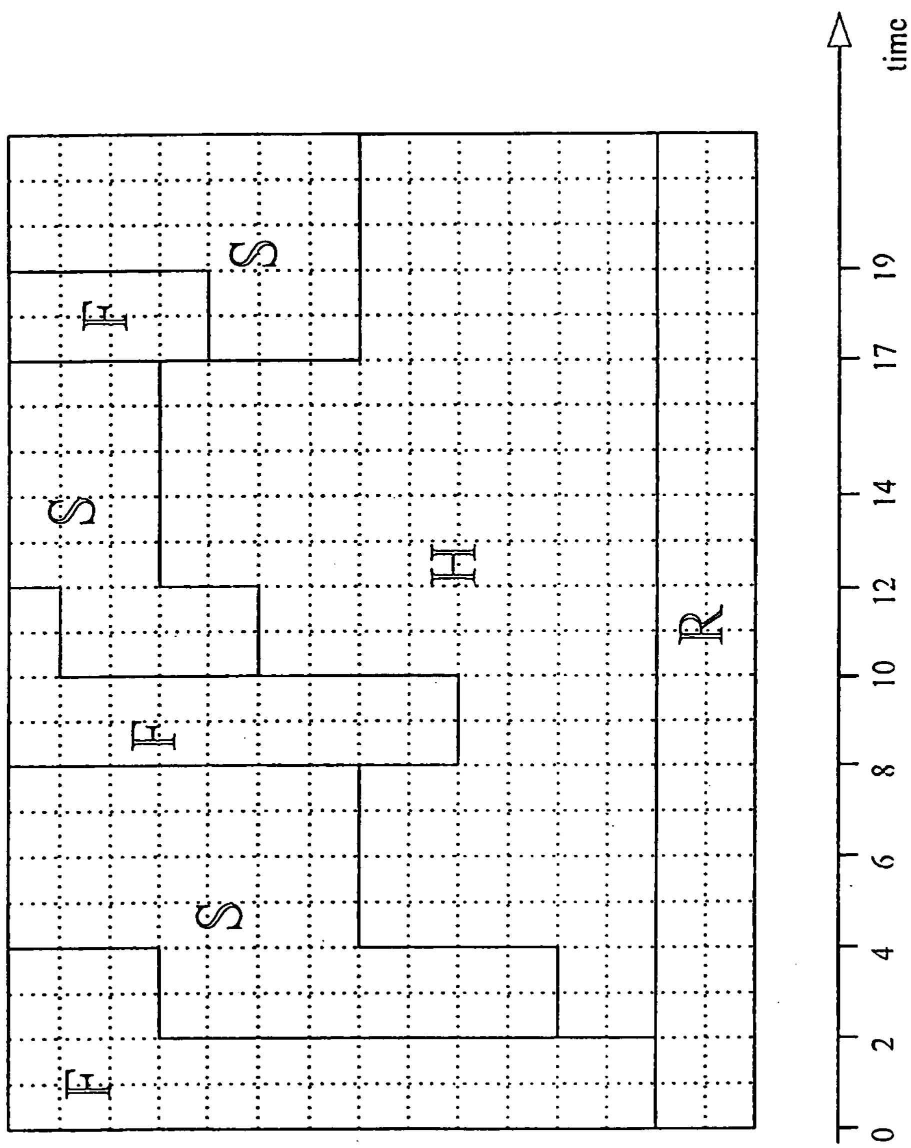

FIG. 4b

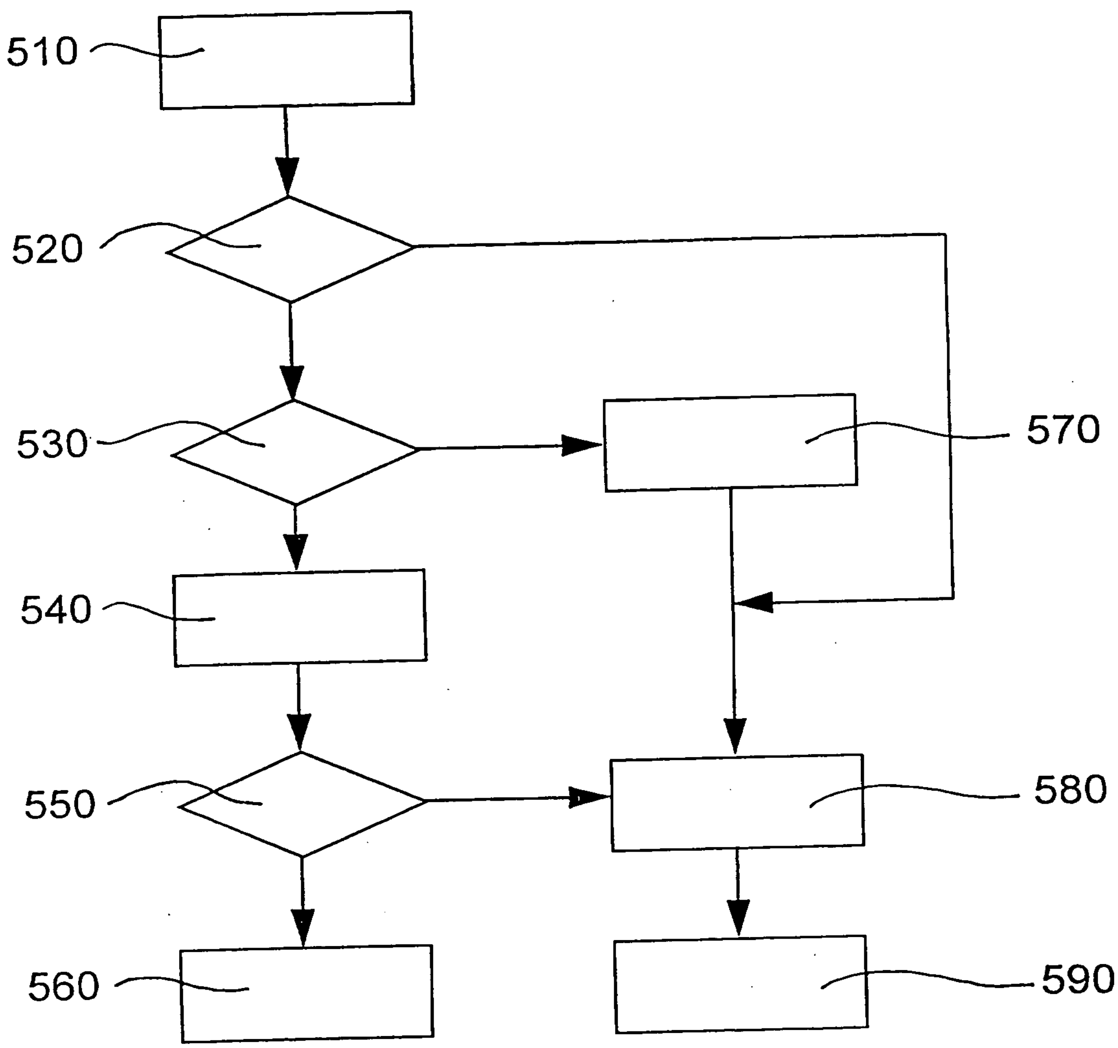

FIG. 5a

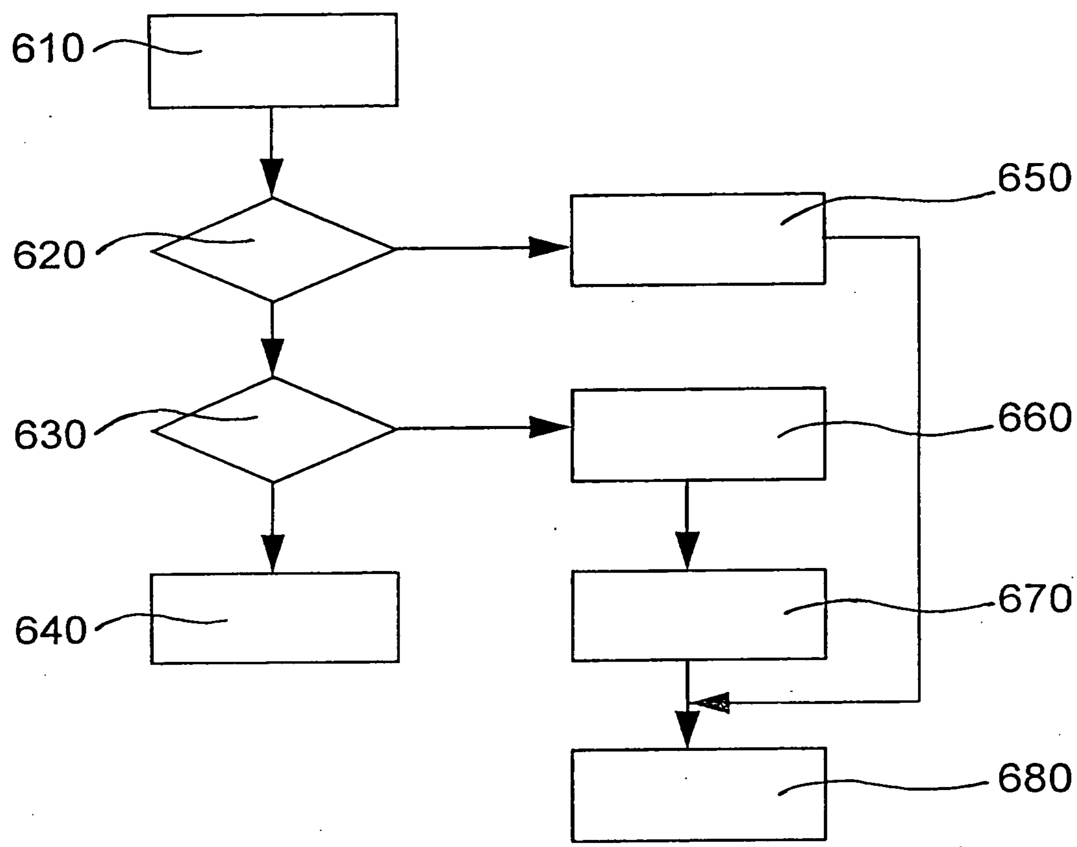

FIG. 5b

slot	channel	priority
3	B	H
4	B	H
5	B	S
6	B	S
7	B	S
8	B	S
9	B	S
10	B	S
11	B	S
12	B	S
13	F	-

FIG. 6a

slot	channel	priority
1	A	R
2	A	R
14	F	-
15	F	-

FIG. 6b

slot	channel	priority
3	B	H
4	B	H
5	B	S
6	B	S
7	B	S
8	B	S
9	B	S
10	B	S
11	B	S
12	C	H
13	C	H
14	C	H
15	C	H

FIG. 6c

slot	channel	priority
1	A	R
2	A	R

FIG. 6d